

MEMORY HUB BYPASS CIRCUIT AND METHOD

TECHNICAL FIELD

5 This invention relates to a computer system, and, more particularly, to a computer system having a memory hub coupling several memory devices to a processor or other memory access device.

BACKGROUND OF THE INVENTION

10 Computer systems use memory devices, such as dynamic random access memory ("DRAM") devices, to store instructions and data that are accessed by a processor. These memory devices are normally used as system memory in a computer system. In a typical computer system, the processor communicates with the system memory through a processor bus and a memory controller. The processor issues a
15 memory request, which includes a memory command, such as a read command, and an address designating the location from which data or instructions are to be read. The memory controller uses the command and address to generate appropriate command signals as well as row and column addresses, which are applied to the system memory. In response to the commands and addresses, data is transferred between the system
20 memory and the processor. The memory controller is often part of a system controller, which also includes bus bridge circuitry for coupling the processor bus to an expansion bus, such as a PCI bus.

Although the operating speed of memory devices has continuously increased, this increase in operating speed has not kept pace with increases in the
25 operating speed of processors. Even slower has been the increase in operating speed of memory controllers coupling processors to memory devices. The relatively slow speed of memory controllers and memory devices limits the data bandwidth between the processor and the memory devices.

In addition to the limited bandwidth between processors and memory devices, the performance of computer systems is also limited by latency problems that increase the time required to read data from system memory devices. More specifically,

5 when a memory device read command is coupled to a system memory device, such as a synchronous DRAM ("SDRAM") device, the read data are output from the SDRAM device only after a delay of several clock periods. Therefore, although SDRAM devices can synchronously output burst data at a high data rate, the delay in initially providing the data can significantly slow the operating speed of a computer system using such SDRAM devices.

10 One approach to alleviating the memory latency problem is to use multiple memory devices coupled to the processor through a memory hub. In a memory hub architecture, a system controller or memory controller is coupled to several memory modules, each of which includes a memory hub coupled to several memory devices. The memory hub efficiently routes memory requests and responses

15 between the controller and the memory devices. Computer systems employing this architecture can have a higher bandwidth because a processor can access one memory device while another memory device is responding to a prior memory access. For example, the processor can output write data to one of the memory devices in the system while another memory device in the system is preparing to provide read data to

20 the processor.

Although computer systems using memory hubs may provide superior performance, they nevertheless often fail to operate at optimum speed for several reasons. For example, even though memory hubs can provide computer systems with a greater memory bandwidth, they still suffer from latency problems of the type described

25 above. More specifically, although the processor may communicate with one memory device while another memory device is preparing to transfer data, it is sometimes necessary to receive data from one memory device before the data from another memory device can be used. In the event data must be received from one memory device before data received from another memory device can be used, the latency

MEMORY HUB BYPASS CIRCUIT AND METHOD

TECHNICAL FIELD

5 This invention relates to a computer system, and, more particularly, to a computer system having a memory hub coupling several memory devices to a processor or other memory access device.

BACKGROUND OF THE INVENTION

10 Computer systems use memory devices, such as dynamic random access memory ("DRAM") devices, to store instructions and data that are accessed by a processor. These memory devices are normally used as system memory in a computer system. In a typical computer system, the processor communicates with the system memory through a processor bus and a memory controller. The processor issues a
15 memory request, which includes a memory command, such as a read command, and an address designating the location from which data or instructions are to be read. The memory controller uses the command and address to generate appropriate command signals as well as row and column addresses, which are applied to the system memory. In response to the commands and addresses, data is transferred between the system
20 memory and the processor. The memory controller is often part of a system controller, which also includes bus bridge circuitry for coupling the processor bus to an expansion bus, such as a PCI bus.

Although the operating speed of memory devices has continuously increased, this increase in operating speed has not kept pace with increases in the
25 operating speed of processors. Even slower has been the increase in operating speed of memory controllers coupling processors to memory devices. The relatively slow speed of memory controllers and memory devices limits the data bandwidth between the processor and the memory devices.

In addition to the limited bandwidth between processors and memory devices, the performance of computer systems is also limited by latency problems that increase the time required to read data from system memory devices. More specifically, when a memory device read command is coupled to a system memory device, such as a 5 synchronous DRAM ("SDRAM") device, the read data are output from the SDRAM device only after a delay of several clock periods. Therefore, although SDRAM devices can synchronously output burst data at a high data rate, the delay in initially providing the data can significantly slow the operating speed of a computer system using such SDRAM devices.

10 One approach to alleviating the memory latency problem is to use multiple memory devices coupled to the processor through a memory hub. In a memory hub architecture, a system controller or memory controller is coupled to several memory modules, each of which includes a memory hub coupled to several memory devices. The memory hub efficiently routes memory requests and responses 15 between the controller and the memory devices. Computer systems employing this architecture can have a higher bandwidth because a processor can access one memory device while another memory device is responding to a prior memory access. For example, the processor can output write data to one of the memory devices in the system while another memory device in the system is preparing to provide read data to 20 the processor.

Although computer systems using memory hubs may provide superior performance, they nevertheless often fail to operate at optimum speed for several reasons. For example, even though memory hubs can provide computer systems with a greater memory bandwidth, they still suffer from latency problems of the type described 25 above. More specifically, although the processor may communicate with one memory device while another memory device is preparing to transfer data, it is sometimes necessary to receive data from one memory device before the data from another memory device can be used. In the event data must be received from one memory device before data received from another memory device can be used, the latency

problem continues to slow the operating speed of such computer systems. In addition, the memory hub is designed to handle multiple memory requests. Thus, it is only when the memory hub is busy servicing more than one memory request that the benefits of communicating with multiple memory requests are actually realized. Thus, when the
5 memory hub is not busy, the slower and more complex logic used by the memory hub to handle multiple memory requests creates additional latency when servicing only one memory request.

There is therefore a need for a memory hub that bypasses the normal logic used to handle multiple memory requests when only one memory request is being
10 serviced.

SUMMARY OF THE INVENTION

The present invention is directed to a computer system and method of accessing a plurality of memory devices with a memory hub. The computer system
15 includes a plurality of memory modules coupled to a memory hub controller. Each of the memory modules includes the plurality of memory devices and the memory hub. The memory hub includes a link interface, a sequencer, a bypass circuit, and a memory device interface. The link interface receives memory requests from the memory hub controller and forwards the memory requests to either the sequencer or both the
20 sequencer and the bypass circuit based on the status of the memory device interface. The memory device interface couples memory requests to the memory devices. When the memory device interface is busy servicing one or more memory requests, the sequencer generates memory requests and couples the memory requests to the memory device interface. When the memory device interface is not busy servicing one or more
25 memory requests, the bypass circuit generates memory requests and couples a portion of each of the memory requests to the memory device interface. The sequencer generates and couples the remaining portion of each of the memory requests to the memory device interface. The bypass circuit allows the memory requests to more quickly access the memory devices when the memory device interface is not busy,

thereby avoiding the additional latency that would otherwise be created by the sequencer.

As will be apparent, the invention is capable of modifications in various obvious aspects, all without departing from the spirit and scope of the present 5 invention. Accordingly, the drawings and detailed description are to be regarded as illustrative in nature and not restrictive.

BRIEF DESCRIPTION OF THE DRAWINGS

Figure 1 is a block diagram of a computer system according to one 10 example of the invention in which a memory hub is included in each of a plurality of memory modules.

Figure 2 is a block diagram of a memory hub used in the computer system of Figure 1.

15 DETAILED DESCRIPTION OF THE INVENTION

A computer system 100 according to one example of the invention is shown in Figure 1. The computer system 100 includes a processor 104 for performing various computing functions, such as executing specific software to perform specific calculations or tasks. The processor 104 includes a processor bus 106 that normally 20 includes an address bus, a control bus, and a data bus. The processor bus 106 is typically coupled to cache memory 108, which, as previously mentioned, is usually static random access memory (“SRAM”). Finally, the processor bus 106 is coupled to a system controller 110, which is also sometimes referred to as a “North Bridge” or “memory controller.”

25 The system controller 110 serves as a communications path to the processor 104 for a variety of other components. More specifically, the system controller 110 includes a graphics port that is typically coupled to a graphics controller 112, which is, in turn, coupled to a video terminal 114. The system controller 110 is also coupled to one or more input devices 118, such as a keyboard or a mouse, to allow

an operator to interface with the computer system 100. Typically, the computer system 100 also includes one or more output devices 120, such as a printer, coupled to the processor 104 through the system controller 110. One or more data storage devices 124 are also typically coupled to the processor 104 through the system controller 110 to 5 allow the processor 104 to store data or retrieve data from internal or external storage media (not shown). Examples of typical storage devices 124 include hard and floppy disks, tape cassettes, and compact disk read-only memories (CD-ROMs).

The system controller 110 includes a memory hub controller 128 that is coupled to several memory modules 130a, 130b, ...130n, which serve as system 10 memory for the computer system 100. The memory modules 130 are preferably coupled to the memory hub controller 128 through a high-speed link 134, which may be an optical or electrical communication path or some other type of communications path. In the event the high-speed link 134 is implemented as an optical communication path, the optical communication path may be in the form of one or more optical fibers, for 15 example. In such case, the memory hub controller 128 and the memory modules will include an optical input/output port or separate input and output ports coupled to the optical communication path.

The memory modules 130 are shown coupled to the memory hub controller 128 in a multi-drop arrangement in which the single high-speed link 134 is 20 coupled to all of the memory modules 130. However, it will be understood that other topologies may also be used, such as a point-to-point coupling arrangement in which a separate high-speed link (not shown) is used to couple each of the memory modules 130 to the memory hub controller 128. A switching topology may also be used in which the memory hub controller 128 is selectively coupled to each of the memory 25 modules 130 through a switch (not shown). Other topologies that may be used will be apparent to one skilled in the art.

Each of the memory modules 130 includes a memory hub 140 for controlling access to 6 memory devices 148, which, in the example illustrated in Figure 1, are synchronous dynamic random access memory ("SDRAM") devices. However, a

fewer or greater number of memory devices 148 may be used, and memory devices other than SDRAM devices may, of course, also be used. The memory hub 140 is coupled to each of the system memory devices 148 through a bus system 150, which normally includes a control bus, an address bus and a data bus.

5 One example of the memory hub 140 of Figure 1 is shown in Figure 2. The memory hub 140 includes a link interface 152 that is coupled to the high-speed link 134. The nature of the link interface 152 will depend upon the characteristics of the high-speed link 134. For example, in the event the high-speed link 134 is implemented using an optical communications path, the link interface 152 will include an optical input/output port and will convert optical signals coupled through the optical communications path into electrical signals. In any case, the link interface 152 preferably includes a buffer, such as a first-in, first-out buffer 154, for receiving and storing memory requests as they are received through the high-speed link 134. The memory requests are stored in the buffer 154 until they can be processed by the memory hub 140.

10

15

When the memory hub 140 is able to process a memory request, one of the memory requests stored in the buffer 154 is transferred to a memory sequencer 160. The memory sequencer 160 converts the memory requests from the format output by the memory hub controller 128 into a memory request having a format that can be used 20 by the memory devices 148. These re-formatted request signals will normally include memory command signals, which are derived from memory commands contained in the memory requests received by the memory hub 140, and row and column address signals, which are derived from an address contained in the memory requests received by the memory hub 140. In the event one of the memory requests is a write memory 25 request, the re-formatted request signals will normally include write data signals which are derived from write data contained in the memory request received by the memory hub 140. For example, where the memory devices 148 are conventional DRAM devices, the memory sequencer 160 will output row address signals, a row address strobe ("RAS") signal, an active high write/active low read signal ("W/R*"), column

address signals and a column address strobe (“CAS”) signal. The re-formatted memory requests are preferably output from the sequencer 160 in the order they will be used by the memory devices 148. However, the sequencer 160 may output the memory requests in a manner that causes one type of request, such as read requests, to be processed
5 before other types of requests, such as write requests.

The sequencer 160 provides a relatively high bandwidth because it allows the memory hub controller 128 to send multiple memory requests to the memory module 130 containing the memory hub 140, even though previously sent memory requests have not yet been serviced. As a result, the memory requests can be sent at a
10 rate that is faster than the rate at which the memory module 130 can service those requests. The sequencer 160 simply formats the signals of one memory request while memory devices are servicing another memory request. In addition, the sequencer 160 may reorder the memory requests, such as placing a series of read requests before previously received write requests, which reduces the memory read latency.

15 The memory sequencer 160 applies the re-formatted memory requests to a memory device interface 166. The nature of the memory device interface 166 will again depend upon the characteristics of the memory devices 148. In any case, the memory device interface 166 preferably includes a buffer, such as a FIFO buffer 168, for receiving and storing one or more memory requests as they are received from the
20 link interface 152. The memory requests are stored in the buffer 168 until they can be processed by the memory devices 148.

The memory requests are described above as being received by the memory hub 140 in a format that is different from the format that the memory requests are applied to the memory devices 148. However, the memory hub controller 128 may
25 instead re-format the memory requests from the processor 104 (Figure 1) to a format that can be used by the memory devices 148. In such case, it is not necessary for the sequencer 160 to re-format the memory requests. Instead, the sequencer 160 simply schedules the re-formatted memory request signals in the order needed for use by the memory devices 148. The memory request signals for one or more memory requests

are then transferred to the memory device interface 166 so they can subsequently be applied to the memory devices 148.

As previously explained, the sequencer 160 can provide a memory bandwidth that is significantly higher than the memory bandwidth of conventional computer systems. Although the sequencer 160 provides this advantage when the memory hub controller 128 is issuing memory commands at a rapid rate, the sequencer 160 does not provide this advantage when the memory hub controller 128 is issuing memory requests to a memory module 130 at a rate that can be serviced by the memory module 130. In fact, the sequencer 160 can actually increase the read latency of the memory module 130 when no unserviced memory requests are queued in the memory hub 140. The increased latency results from the need to store the memory requests in the sequencer 160, re-format the memory requests, schedule resulting control signals in the sequencer 160, and begin applying those control signals to the memory devices 148. Also, the memory sequencer 160 has a relatively slow clocking structure that can delay the memory hub 140 from issuing to the memory devices memory requests received from the memory hub controller 128.

The memory hub 140 shown in Figure 2 avoids the potential disadvantage of using the memory sequencer 160 by including the bypass circuit 170. The bypass circuit 170 allows the memory requests to access the memory devices 148 more quickly when the memory device interface 166 is not busy servicing at least one memory request. As explained above, when multiple memory requests are not being handled by the sequencer 160, the advantages of servicing memory requests with the sequencer 160 no longer exist. Instead, the sequencer 160 increases the memory read latency. The bypass circuit 170, however, allows the memory hub 140 to decrease the access time of each memory request by handling an initial portion of the signal sequencing normally handled by the sequencer 160, and it preferably uses a faster clocking structure than the sequencer 160. Thus, the bypass circuit 170 increases the access time of the memory requests to the memory devices 148.

The bypass circuit 170 includes conventional circuitry that converts each of the memory requests from the format output by the memory hub controller 128 into a memory request with a format that can be used by the memory devices 148. While the bypass circuit 170 may handle reformatting of the entire memory request, the bypass 5 circuit 170 preferably handles the row address portion of the memory request. Similar to the memory sequencer 160 described above, the bypass circuit 170 receives the memory request from the link interface 154. The bypass circuit 170 then reformats the address portion of the memory request into a row address signal. The bypass circuit 170 outputs the row address signal to the memory device interface 166 and then outputs 10 a row address strobe (RAS) to the memory device interface 166. These signals allow the memory device interface 166 to access the addressed row of one of the memory devices 148. By the time the memory devices have processed the portion of the memory request provided by the bypass circuit 170, the sequencer 160 is ready to provide the remaining portion of the memory request.

15 As shown in Figure 2, the bypass circuit 170 utilizes a link-in clock 176 from the memory hub controller 128 to forward the row address and RAS signals to the memory device interface 166. The link-in clock 176 is received by the link interface 152 and forwarded to the bypass circuit 170. The bypass circuit includes logic that delays and balances the link-in clock 176 with the clock forwarded from the link 20 interface 152 with the memory requests. More specifically, the link-in clock 176 is used to forward each memory request from the memory hub controller 128 to the memory hub 140, in particular to the link interface 152. The link-in clock 176 is then forwarded to the bypass circuit 170. The memory request output by the link interface 152 to the bypass circuit 170 uses a controller clock, which is a slower clock used by 25 the memory hub 140 to process memory requests. The bypass circuit 170 delays and balances the link-in clock 176 with the controller clock, which allows the bypass circuit 170 to use the link-in clock 176 to service the row portion of the memory request. The faster link-in clock 176 allows the bypass circuit 170 to process and forward the row

address and RAS signals more quickly than the controller clock used by the sequencer 160.

While the bypass circuit 170 handles the row portion of the memory request from the link interface 152, the remaining portion of the memory request, for example the command signal and column address, is formatted and forwarded by the sequencer 160. This allows the sequencer 160 to format the remaining portion of the memory request, as explained above, while the read address and RAS signals are accessing the addressed row of one of the memory devices 148. Thus, the sequencer 160 does not have to service the row portion of the memory request. This structure increases the overall access time to the memory devices 148, thus reducing the latency of the memory hub 140, because the bypass circuit 170 forwards the row address and RAS signals to one of the memory devices more quickly than the sequencer 160. In addition, during the clock delays used by the row address and RAS signals to access one of the memory devices 148, the sequencer 160 is formatting and ordering the remaining signals of the memory request. Thus, once the remaining signals are formatted and ordered by the sequencer 160, they can be immediately coupled to the memory device 148 that has already been accessed by the row address signal.

The bypass circuit 170 is utilized by the memory hub 140 when the memory device interface 166 is not busy servicing memory requests. The memory device interface 166 generates a high "ACTIVE/IDLE*" signal when the buffer 168 of the memory device interface 166 is active and contains, for example, one or more memory requests. The high ACTIVE/IDLE* signal indicates that the memory device interface is busy, thus memory requests can be more efficiently handled by using the sequencer 160. When the buffer 168 contains, for example, less than one memory request, the memory device interface generates a low "ACTIVE/IDLE*" signal. The low ACTIVE/IDLE* signal indicates that the memory device interface is not busy, thus the memory hub 140 uses the bypass circuit 170 and sequencer 160 to service memory requests. The ACTIVE and IDLE* conditions generated by the memory device interface 166 are not limited to the circumstances described above. For example, the

memory device interface 166 may generate an ACTIVE signal based on the buffer 168 containing a certain percentage of memory requests and likewise an IDLE* signal when the number of memory requests is under a certain percentage.

The memory hub 140, shown in Figure 2, further includes a multiplexer 172, which works in conjunction with the memory device interface 166 to service the memory requests. The multiplexer 172 has inputs coupled to the bypass circuit 170 and the sequencer 160, an output coupled to the memory device interface 166, and a control input coupled to the memory device interface 166. The multiplexer 172 uses the ACTIVE/IDLE* signal from the memory device interface 166 to couple memory requests to the memory device interface 166. When the multiplexer 172 receives an ACTIVE signal, or a high ACTIVE/IDLE* signal, the multiplexer 172 couples memory requests from the sequencer 160 to the memory device interface 166. Likewise, when the multiplexer 172 receives an IDLE* signal, or a low ACTIVE/IDLE* signal, the multiplexer 172 couples a portion of each memory request from the bypass circuit 170 to the memory device interface 166 and a portion of each memory request from the sequencer 160 to the memory device interface 166.

The ACTIVE/IDLE* signal generated by the memory device interface 166 is also used to determine whether memory requests should be forwarded from the link interface 152 to the sequencer 160 or to both the bypass circuit 170 and the sequencer 160. Both the sequencer 160 and the bypass circuit 170 are coupled to the memory device interface 166. When the memory device interface 166 generates an ACTIVE signal, the sequencer 160 receives the memory requests from the link interface 152 and generates and couples memory requests to the multiplexer 172. When the memory device interface 166 generates an IDLE* signal, both the sequencer 160 and the bypass circuits receive the memory requests and handle specific portions of each of the memory requests, as described above.

Although the present invention has been described with reference to preferred embodiments, persons skilled in the art will recognize that changes may be

made in form and detail without departing from the spirit and scope of the invention.
Accordingly, the invention is not limited except as by the appended claims.

CLAIMS

1. A memory module, comprising:
 - a plurality of memory devices; and
 - a memory hub, comprising:
 - a link interface receiving memory requests for access to at least one of the memory devices;

a memory device interface coupled to the memory devices, the memory device interface coupling memory requests to the memory devices and generating a status signal indicating whether or not at least one of the memory requests is being serviced;

a bypass circuit coupled to the link interface and the memory device interface, the bypass circuit generating and coupling a portion of each of the memory requests from the link interface to the memory device interface responsive to the status signal from the memory device interface indicating that at least one of the memory requests is not being serviced; and

a sequencer coupled to the link interface and the memory device interface, the sequencer generating and coupling memory requests from the link interface to the memory device interface responsive to the status signal from the memory device interface indicating that at least one of the memory requests is being serviced and generating and coupling the remaining portion of each of the memory requests not handled by the bypass circuit from the link interface to the memory device interface responsive to the status signal from the memory device interface indicating that at least one of the memory requests is not being serviced.

2. The memory module of claim 1 wherein the status signal comprises:

an active signal indicating that at least one of the memory requests is being serviced; and

an idle signal indicating that at least one of the memory requests is not being serviced.

3. The memory module of claim 1 wherein the memory hub further comprises a multiplexer having data inputs coupled to the sequencer and the bypass circuit, a data output coupled to the memory device interface and a control input coupled to receive the status signal from the memory device interface, the multiplexer coupling the memory requests from the sequencer responsive to the status signal indicating that at least one of the memory requests is being serviced and coupling a portion of the memory request from the bypass circuit and a portion from the sequencer responsive to the status signal indicating that at least one of the memory requests is not being serviced.

4. The memory module of claim 1 wherein the link interface further receives a link-in clock with each of the memory requests, the link-in clock being forwarded to the bypass circuit along with the memory requests and used by the bypass circuit to generate and couple memory requests to the memory device interface responsive to the status signal from the memory device interface indicating that at least one of the memory requests is not being serviced.

5. The memory module of claim 3 wherein the link interface further receives a link-in clock with each of the memory requests, the link-in clock being forwarded to the bypass circuit along with the memory requests and used by the bypass circuit to generate and couple memory requests to the multiplexer responsive to the status signal indicating that at least one of the memory requests is not being serviced.

6. The memory module of claim 1 wherein the memory device interface comprises a first-in, first-out buffer that is operable to receive and store

memory requests received from the link interface and to transfer the stored memory requests to at least one of the memory devices in the order in which they were received.

7. The memory module of claim 1 wherein the link interface comprises a first-in, first-out buffer that is operable to receive and store memory requests and to transfer the stored memory requests to the memory device interface in the order in which they were received.

8. The memory module of claim 6 wherein the link interface comprises a first-in, first-out buffer that is operable to receive and store memory requests and to transfer the stored memory requests to the memory device interface in the order in which they were received.

9. The memory module of claim 1 wherein the memory device interface is coupled to the link interface, the memory device interface further receiving read data responsive to the memory requests and coupling the read data to the link interface.

10. The memory module of claim 1 wherein the link interface comprises an optical input/output port.

11. The memory module of claim 1 wherein the memory devices comprises dynamic random access memory devices.

12. The memory module of claim 1 wherein the portion of each of the memory requests generated and coupled by the bypass circuit includes the row address portion of each of the memory requests.

13. A memory hub, comprising:

a link interface receiving memory requests;

a memory device interface operable to output memory requests, the memory device interface generating a status signal indicating whether or not at least one of the memory requests is being serviced;

a bypass circuit coupled to the link interface and the memory device interface, the bypass circuit generating and coupling a portion of each of the memory requests from the link interface to the memory device interface responsive to the status signal from the memory device interface indicating that at least one of the memory requests is not being serviced; and

a sequencer coupled to the link interface and the memory device interface, the sequencer generating and coupling the memory requests from the link interface to the memory device interface responsive to the status signal from the memory device interface indicating that at least one of the memory requests is being serviced and generating and coupling the remaining portion of each of the memory requests not handled by the bypass circuit from the link interface to the memory device interface responsive to the status signal from the memory device interface indicating that at least one of the memory requests is not being serviced.

14. The memory hub of claim 12 wherein the status signal comprises:

an active signal indicating that at least one of the memory requests is being serviced; and

an idle signal indicating that at least one of the memory requests is not being serviced.

15. The memory hub of claim 13 wherein the memory hub further comprises a multiplexer having data inputs coupled to the sequencer and the bypass circuit, a data output coupled to the memory device interface and a control input

coupled to receive the status signal from the memory device interface, the multiplexer coupling the memory requests from the sequencer responsive to the status signal indicating that at least one of the memory requests is being serviced and coupling a portion of the memory request from the bypass circuit and a portion from the sequencer responsive to the status signal indicating that at least one of the memory requests is not being serviced.

16. The memory hub of claim 13 wherein the link interface further receives a link-in clock with each of the memory requests, the link-in clock being forwarded to the bypass circuit along with the memory requests and used by the bypass circuit to generate and couple memory requests to the memory device interface responsive to the status signal from the memory device interface indicating that at least one of the memory requests is not being serviced.

17. The memory hub of claim 15 wherein the link interface further receives a link-in clock with each of the memory requests, the link-in clock being forwarded to the bypass circuit along with the memory requests and used by the bypass circuit to generate and couple memory requests to the multiplexer responsive to the status signal indicating that at least one of the memory requests is not being serviced.

18. The memory hub of claim 13 wherein the memory device interface comprises a first-in, first-out buffer that is operable to receive and store memory requests received from the link interface and to transfer the stored memory requests to at least one of the memory devices in the order in which they were received.

19. The memory hub of claim 13 wherein the link interface comprises a first-in, first-out buffer that is operable to receive and store memory requests and to transfer the stored memory requests to the memory device interface in the order in which they were received.

20. The memory hub of claim 18 wherein the link interface comprises a first-in, first-out buffer that is operable to receive and store memory requests and to transfer the stored memory requests to the memory device interface in the order in which they were received.

21. The memory hub of claim 13 wherein the memory device interface is coupled to the link interface, the memory device interface further receiving read data responsive to the memory requests and coupling the read data to the link interface.

22. The memory hub of claim 13 wherein the link interface comprises an optical input/output port.

23. The memory hub of claim 13 wherein the memory devices comprises dynamic random access memory devices.

24. The memory hub of claim 13 wherein the portion of each of the memory requests generated and coupled by the bypass circuit includes the row address portion of each of the memory requests.

25. A computer system, comprising:
a central processing unit ("CPU");
a system controller coupled to the CPU, the system controller having an input port and an output port;
an input device coupled to the CPU through the system controller;
an output device coupled to the CPU through the system controller;
a storage device coupled to the CPU through the system controller; and

a plurality of memory modules, each of the memory modules comprising:

a plurality of memory devices; and

a memory hub, comprising:

a link interface receiving memory requests for access to at least one of the memory devices;

a memory device interface coupled to the memory devices, the memory device interface coupling memory requests to the memory devices and generating a status signal indicating whether or not at least one of the memory requests is being serviced;

a bypass circuit coupled to the link interface and the memory device interface, the bypass circuit generating and coupling a portion of each of the memory requests from the link interface to the memory device interface responsive to the status signal from the memory device interface indicating that at least one of the memory requests is not being serviced; and

a sequencer coupled to the link interface and the memory device interface, the sequencer generating and coupling memory requests from the link interface to the memory device interface responsive to the status signal from the memory device interface indicating that at least one of the memory requests is being serviced and generating and coupling the remaining portion of each of the memory requests not handled by the bypass circuit from the link interface to the memory device interface responsive to the status signal from the memory device interface indicating that at least one of the memory requests is not being serviced; and

a communications link coupling the output port of the system controller to the input port of the memory hub in each of the memory modules, and coupling the input port of the system controller to the output port of the memory hub in each of the memory modules.

26. The computer system of claim 26 wherein the status signal comprises:

an active signal indicating that at least one of the memory requests is being serviced; and

an idle signal indicating that at least one of the memory requests is not being serviced.

27. The computer system of claim 26 wherein the memory hub further comprises a multiplexer having data inputs coupled to the sequencer and the bypass circuit, a data output coupled to the memory device interface and a control input coupled to receive the status signal from the memory device interface, the multiplexer coupling the memory requests from the sequencer responsive to the status signal indicating that at least one of the memory requests is being serviced and coupling a portion of the memory request from the bypass circuit and a portion from the sequencer responsive to the status signal indicating that at least one of the memory requests is not being serviced.

28. The computer system of claim 26 wherein the link interface further receives a link-in clock with each of the memory requests, the link-in clock being forwarded to the bypass circuit along with the memory requests and used by the bypass circuit to generate and couple memory requests to the memory device interface responsive to the status signal from the memory device interface indicating that at least one of the memory requests is not being serviced.

29. The computer system of claim 27 wherein the link interface further receives a link-in clock with each of the memory requests, the link-in clock being forwarded to the bypass circuit along with the memory requests and used by the bypass circuit to generate and couple memory requests to the multiplexer responsive to

the status signal indicating that at least one of the memory requests is not being serviced.

30. The computer system of claim 26 wherein the memory device interface comprises a first-in, first-out buffer that is operable to receive and store memory requests received from the link interface and to transfer the stored memory requests to at least one of the memory devices in the order in which they were received.

31. The computer system of claim 26 wherein the link interface comprises a first-in, first-out buffer that is operable to receive and store memory requests and to transfer the stored memory requests to the memory device interface in the order in which they were received.

32. The computer system of claim 30 wherein the link interface comprises a first-in, first-out buffer that is operable to receive and store memory requests and to transfer the stored memory requests to the memory device interface in the order in which they were received.

33. The computer system of claim 26 wherein the memory device interface is coupled to the link interface, the memory device interface further receiving read data responsive to the memory requests and coupling the read data to the link interface.

34. The computer system of claim 26 wherein the link interface comprises an optical input/output port.

35. The computer system of claim 26 wherein the memory devices comprises dynamic random access memory devices.

36. The computer system of claim 26 wherein the portion of each of the memory requests generated and coupled by the bypass circuit includes the row address portion of each of the memory requests.

37. The computer system of claim 26 wherein the input and output ports of the system controller comprises a combined input/output port coupled to the communications link, and wherein the input and output ports of each of the memory hubs comprises a combined input/output port coupled to the communications link.

38. The computer system of claim 34 wherein the communications link comprises an optical communications link, wherein the input and output ports of the system controller comprises an optical input/output port coupled to the optical communications link and wherein the input and output ports of each of the memory hubs comprises a respective optical input/output port coupled to the optical communications link.

39. A method of accessing data in each of a plurality of memory devices on each of a plurality of memory modules, each of the memory modules including a memory hub, the method comprising:

checking if a memory device interface located on the memory hub is servicing memory requests;

if the memory device interface is servicing memory requests, sending memory requests through a sequencer located on the memory hub to the memory device interface; and

if the memory device interface is not busy servicing memory requests, sending a portion of each of the memory requests through a bypass circuit located on the memory hub to the memory device interface and the remaining portion of each of the memory requests through the sequencer to the memory device interface.

40. The method of claim 39 wherein the portion of the each of the memory requests sent through the bypass circuit includes the row address portion of each of the memory requests.

MEMORY HUB BYPASS CIRCUIT AND METHOD

ABSTRACT OF THE DISCLOSURE

A computer system and a method used to access data from a plurality of memory devices with a memory hub. The computer system includes a plurality of memory modules coupled to a memory hub controller. Each of the memory modules includes the memory hub and the plurality of memory devices. The memory hub includes a sequencer and a bypass circuit. When the memory hub is busy servicing one or more memory requests, the sequencer generates and couples the memory requests to the memory devices. When the memory hub is not busy servicing multiple memory requests, the bypass circuit generates and couples a portion of each the memory requests to the memory devices and the sequencer generates and couples the remaining portion of each of the memory requests to the memory devices.

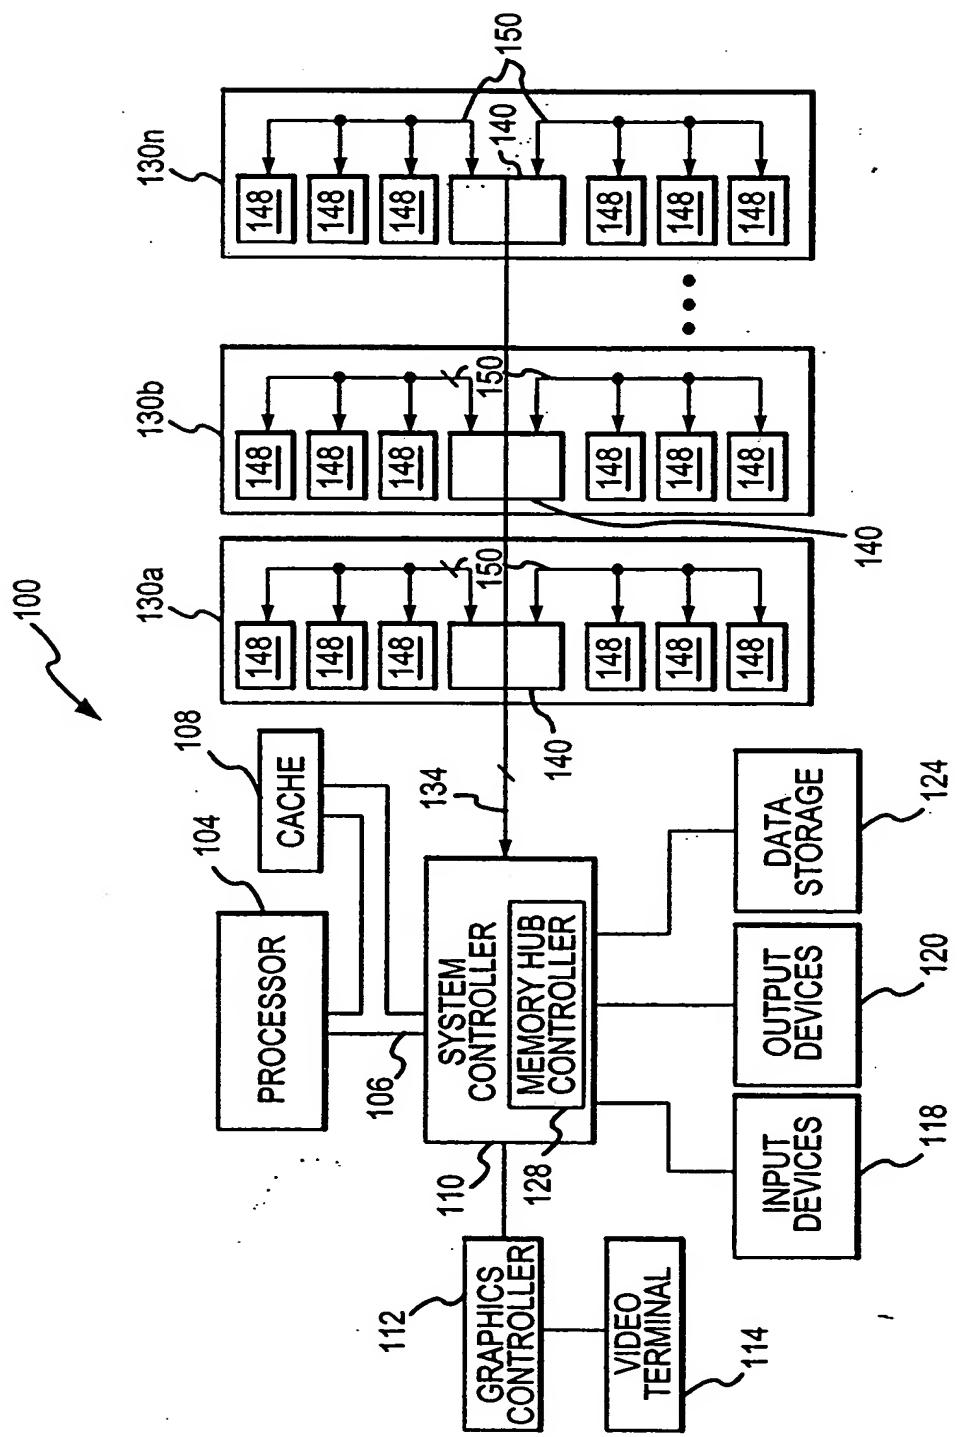


FIG. 1

FIG.2

